ELSEVIER

Journal of Archaeological Science 31 (2004) 533–549

Journal of Archaeological SCIENCE

http://www.elsevier.com/locate/jas

Comparing archaeological settlement systems with rank-size graphs: a measure of shape and statistical confidence

Robert D. Drennan*, Christian E. Peterson

Department of Anthropology, University of Pittsburgh, 3H01 Wesley W Posvar Hall, Pittsburgh, PA 15260, USA

Received 11 July 2003; received in revised form 22 September 2003; accepted 6 October 2003

Abstract

Rank-size analysis of settlement systems in archaeology has focused strongly on departures from log-normality, but it is not clear that a log-normal pattern should ordinarily be expected. Direct comparison of observed rank-size curves to each other is of greater utility in identifying chronological change and inter-regional variation in settlement dynamics as reflected in rank-size patterns. A coefficient of shape is proposed to facilitate comparison and to serve as a basis for bootstrap assessment of statistical confidence in any differences that are observed. The approach is illustrated with analysis of archaeological settlement data from China, Iran, Peru, and Mexico. These analyses provide insight into how many settlements must be observed to achieve high confidence and into the impact of curve shape on confidence.
© 2003 Elsevier Ltd. All rights reserved.

Keywords: Rank-size analysis; Significance testing; Settlement patterns; Yellow River Valley; Susiana Plain; Santa Valley; Valley of Oaxaca

1. Introduction

The rank-size graph was adopted into regional settlement studies in archaeology from geography nearly 30 years ago [2,6,9–11,13,18–20]. For geographers it has been a way of characterizing the evenness or unevenness of population distribution across the settlements in a region. In any region, there will be some variation in the sizes of different settlements, and it is of interest to consider the extent to which a regional population is concentrated into the single largest town or city, as opposed to being spread more evenly across a number of settlements of more nearly equal size. The rank-size graph is a tool for characterizing this aspect of population distribution. The relationship between settlement rank (with the largest settlement being rank 1, the next largest rank 2, etc.) and settlement size has typically been graphed as a line plot of the logarithms of both variables. Since logarithms are used, a few very high-ranking settlements, which are usually those of greatest interest, have a very large impact on the overall shape of the line on the graph, while many low-ranking settlements are crowded into the lower section of the line.

The *rank-size rule* suggests that we might "expect" the rank 2 settlement to be half as large as the rank 1 settlement; the rank 3 settlement to be one-third as large as the rank 1 settlement; and so on [25]. If such a relationship is subjected to a log transformation, the result is a straight line on a rank-size graph, running from upper left to lower right. This straight-line pattern has been dubbed *log-normal* in the rank-size literature, and, because it is a straight line, it is easy to spot deviations from it. Other patterns are often described in terms of how they depart from the log-normal line: the overwhelmingly large population of the largest settlement in a *primate* distribution makes the curve drop rapidly below the log-normal line, whereas several large settlements with similar populations result in a *convex* pattern (Fig. 1). Several variations on this theme created by combining these departures from log-normal in different ways have also been identified e.g., [21]. The generality of the assumptions underlying the log-normal expectation has been questioned in contemporary

* Corresponding author. Tel.: +1-412-648-7500; fax: +1-412-648-7535
E-mail address: drennan@pitt.edu (R.D. Drennan).

0305-4403/04/$ - see front matter © 2003 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jas.2003.10.002

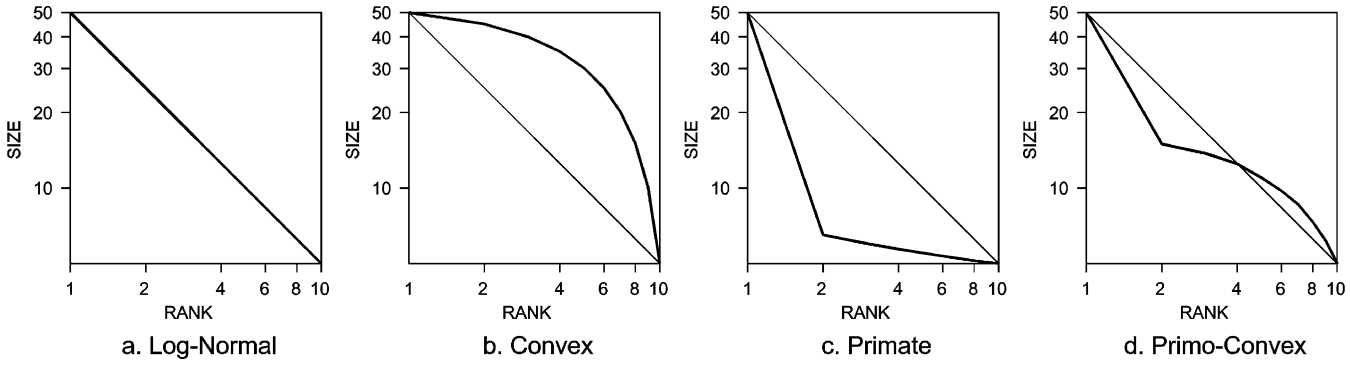

Fig. 1. Examples of rank-size graphs with different shapes.

studies, and seems even more suspect for the range of societies whose settlement distributions have been studied in this way archaeologically [11 (p. 148), 19 (pp. 261–267), 20]. It may be of greater utility, especially in archaeological studies, to compare rank-size plots of settlement distributions (of different periods or in different regions) to each other than to focus on the extent to which any one of them departs from the log-normal expectation. In this paper we propose analytical tools to assist in making such comparisons and to assess statistical confidence in them.

2. Characterizing the shape of rank-size curves: the coefficient *A*

Archaeologists have become accustomed to subjective characterizations of the shape of rank-size curves, using labels such as those in Fig. 1. The primate curve in Fig. 1c is clearly quite different from the log-normal pattern (Fig. 1a), but primate curves differ from one another as well. In comparing one primate curve to another, one might like to improve upon descriptions such as "A is a primate pattern, but B is even more primate than A." A mathematical coefficient for describing the shape of a rank-size curve would be helpful for putting a finer point on such comparisons, and could serve as a foundation for establishing levels of confidence. A coefficient based on the measurement of areas on opposite sides of the log-normal line would relate easily and intuitively to the way in which we are accustomed to visualizing and describing rank-size patterns.

The log-normal line in a rank-size graph, running from the largest settlement at the upper left to the smallest at the lower right, defines a rectangular frame. The coefficient of shape proposed here takes the area of this rectangle as its basis. The total area inside the rectangular frame is, of course, divided into two equal parts by the log-normal line, because the log-normal line is the diagonal of the rectangle. Add an observed rank-size curve, and either or both of these two halves may be further subdivided. A convex pattern, as in Fig. 2a for example, can be characterized mathematically as the shaded area between the log-normal line and the observed rank-size curve. The more strongly convex pattern in Fig. 2b would be characterized by a number roughly twice as large, representing the larger shaded area. In a similar fashion, a concave, or primate, pattern can be characterized by the area between itself and the log-normal line. If this area is given a negative value, since it is below the log-normal line rather than above it, a primate pattern is distinguished from a convex pattern by its sign. The primate pattern in Fig. 2d has a negative area roughly twice as large as that in Fig. 2c, and thus a coefficient roughly twice as far below zero as the one for Fig. 2c. The analogous area for a perfect log-normal pattern is equal to zero, since, of course, there is no area at all between it and the log-normal line. The coefficient of shape proposed here, *A*, is equal to the area between the observed rank-size pattern and the log-normal line. Areas measured above the log-normal line are considered positive; and those below it, negative. Consequently, a rank-size pattern that crosses from one side of the log-normal line to the other, yields both positive and negative components, which are summed to yield the value of the *A* coefficient. In such a case, if the area above the log-normal line is larger than the area below it, the value of *A* is positive (Fig. 2e). If the area below the log-normal line is larger than the area above it, the value of *A* is negative (Fig. 2f).

If the rank-size graph is scaled so that the rectangle defined by the endpoints of the log-normal line is square with an area equal to 2.0 units (1.0 unit above the log-normal line and 1.0 below it), then the *A* coefficient has a usefully delimited range. In theory, *A* can attain a maximum value of 1.0. Such a value represents a complete lack of settlement hierarchy, since it only occurs when all settlements have exactly the same size and the rank-size graph is a horizontal line across the top of the frame (all sites tied for rank 1). In practice a value of -1.0 for *A* indicates the opposite pattern: extreme primateness, the overwhelming dominance of the settlement hierarchy by a single very large settlement. This is

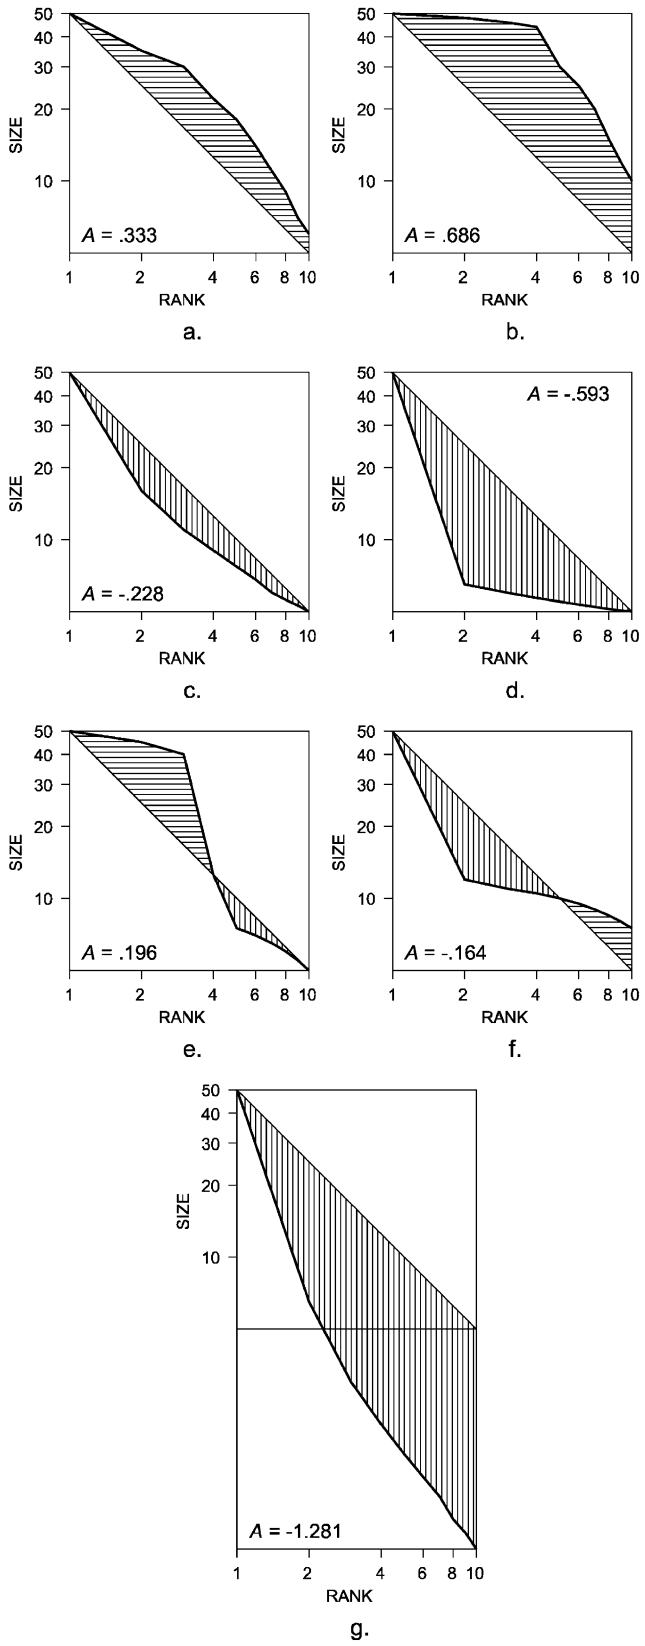

Fig. 2. Areas in a rank-size graph used as positive and negative components of the coefficient A (horizontal and vertical shading, respectively).

not, however, a theoretical minimum, as values even lower can be produced if one or more observed settlement sizes are less than the "expected" size of the smallest settlement in the ideal log-normal pattern (Fig. 2g). For example, in a sample of 15 settlements, the "expected" size of the smallest settlement, according to the rank-size rule, would be 1/15 the size of the largest settlement. It is certainly possible that one or more actual settlements could be smaller than this, and if many were (or if some were much smaller than this) the value of A could be even less than -1.0. A is an elaboration of a useful, though little-used, rank-size convexity index proposed by Johnson [10 (pp. 237–239), 11 (pp. 154–155)]. While Johnson's index gives equal weight to all settlements, the A coefficient gives greater weight to the higher ranking settlements, as does the rank-size graph.

The coefficient A, then, provides a numerical value close to 0.0 for patterns roughly similar to log-normal (either following the log-normal line closely or delineating roughly equal positive and negative areas). Positive values, ranging up to a theoretical maximum of 1.0, indicate increasingly convex patterns, whereas negative values, ranging down to or even beyond -1.0, indicate increasingly concave, or primate, patterns. Since A is really a measure of the net tendency of a rank-size curve toward convexity or concavity, curves of somewhat different shapes can produce similar A values. Consequently, it is important to use the A coefficient as a complement to (not a substitute for) actually examining the rank-size graph. Figs. 1 and 2 illustrate the shapes associated with varying values of A. The A scale, roughly centered on 0.0, automatically indicates the strength of the departure of an observed pattern from log-normal, and it facilitates direct comparison between two or more observed patterns (as in two different periods or regions). Such comparisons, like most, are meaningful only if the differences observed have sufficient statistical significance. For example, a change in rank-size pattern from one period to another may be taken to indicate some change in regional political or economic centralization. Before attempting to account for such a political or economic change, however, it is essential to assess the probability that a difference in rank-size pattern of the magnitude observed could be the result of nothing more than the vagaries of sampling. The A coefficient provides a basis for such an assessment, by establishing an error range for the level of statistical confidence desired.

3. Monte Carlo sampling

The statistical significance of an observed pattern depends on the strength of the result and the size of the sample in which it is observed. For rank-size patterns, like other parameters, the significance of the difference

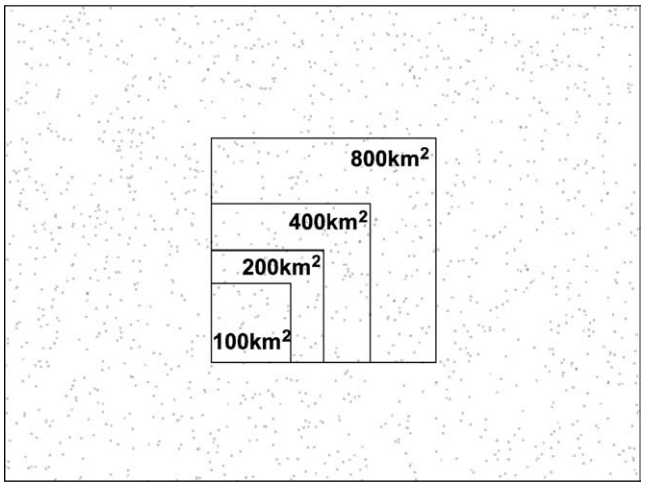

Fig. 3. Randomly generated population of 1452 settlements, with sample blocks of the four different sizes employed in the Monte Carlo experiment.

between two observed patterns (or between one observed pattern and log-normal) will be greater if the difference is greater and/or if the number of observations is larger. A simple Monte Carlo experiment enables us to explore rank-size patterns in populations of settlements and in samples drawn from them. This exploration proves to be enlightening primarily for what it shows about why Monte Carlo sampling is *not* a useful approach to significance testing for rank-size graphs. We begin with a randomly generated population of 1452 settlements (Fig. 3). Settlement areas correspond to a Poisson distribution, with a mean of 3.2 ha; the largest covers 15.1 ha; the four smallest are less than 0.01 ha. The X and Y coordinates of their locations were randomly generated within a 60 by 80 km area.

Samples of four different sizes were selected from this population by randomly locating squares of 100 km^2, 200 km^2, 400 km^2, and 800 km^2 (Fig. 3). The Monte Carlo experiment thus resembles full-coverage survey of a contiguous tract imagined as a sample from a larger region. For each block size, 1000 different samples were selected by randomly generating coordinates for southwest corners of sample blocks. Values for A were calculated for each sample survey block; their means and standard deviations are presented in Table 1. At first glance, this would seem an appropriate model for thinking about sampling and statistical confidence issues for rank-size patterns. If the samples from this Monte Carlo experiment were taken to represent the sampling distributions one could expect for such survey under parametric statistical assumptions, these standard deviations would represent error ranges at about the 66% confidence level. We would thus appear to be in position to estimate the A value for the entire population as 0.537 ± 0.111 from the smallest sample block with about a two-thirds chance that the "true" population value

Table 1
Means and standard deviations for A values for sample blocks of different sizes from the randomly generated population of 1452 settlements

Sample block size (km^2)	Number of samples	Mean A	Standard deviation
100	1000	0.537	0.111
200	1000	0.621	0.071
400	1000	0.686	0.050
800	1000	0.734	0.035

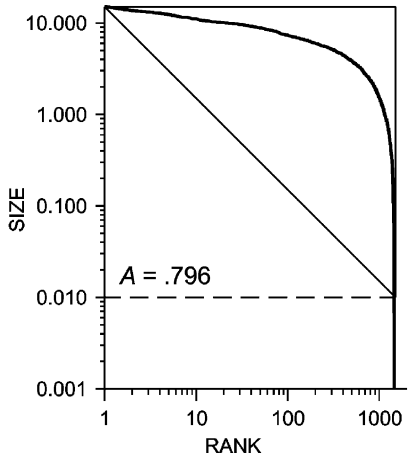

Fig. 4. Rank-size graph for randomly generated population of 1452 settlements.

would lie within this interval. Much as we would expect, the standard deviation decreases as the size of the sample block increases. Surprisingly, however, the average estimate of the A value also changes sharply as the size of the sample block increases. This is distressing, since it would lead us to make very different estimates of the A value for the population depending on the sample block size chosen, and these differences in the A values are considerably larger than the standard deviations we might be tempted to use as error ranges. Furthermore, none of the estimates would be very good, since the actual value of A for the entire population of 1452 settlements is 0.796. This indicates a substantially more convex rank-size pattern (Fig. 4) than any of the samples would, one outside the 66% confidence level error range for even the largest sample.

Moreover, examination of the frequency distributions of the 1000 A values for sample blocks of each size reveals that the shapes of all four distributions depart markedly from the normal shape that would permit us to use the standard deviations in the way we just have (Fig. 5). The two smaller sample blocks yield asymmetrical distributions, and the larger two are multimodal. This clearly contributes to the inadequacy of using the standard deviation as an error range for 66% confidence,

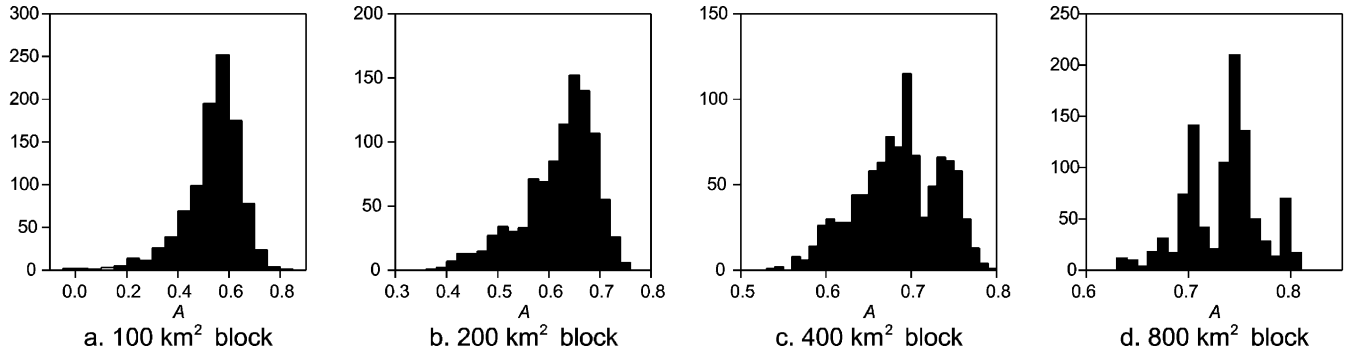

Fig. 5. Frequency distributions of A values from four different sample block sizes in the Monte Carlo experiment with the randomly generated population of 1452 settlements.

but a consideration of the factors that produce these intractable shapes leads to an even more important point. One or a few large settlements, of course, have a major impact on the shape of a rank-size curve. It appears that, in the Monte Carlo experiment, the multimodality observed in frequency distributions of A for larger sample blocks reflects several different kinds of results: sample blocks that capture one especially large settlement, those that capture two especially large settlements, and so on. These different kinds of results yield A values that cluster around different modes, producing the multimodal shapes seen in Fig. 5c and d and contributing to the asymmetries of Fig. 5a and b. The different sample block sizes, then, produce not just larger or smaller samples of settlements, but in effect represent spatial patterning at fundamentally different scales. At the scale of the 100 km^2 block, the rank-size pattern is least convex. At larger scales, it is increasingly common for sample blocks to capture several of the largest settlements, producing a much more convex pattern. And at the scale of the entire population, the inclusion of all the larger settlements produces the most convex curve of all. The differences in A values, then, between the analytical blocks of different sizes is not just a peculiarity of the A coefficient; it is a real difference of some magnitude between rank-size patterns, which the A coefficient calls to our attention and provides a useful index of.

Since this effect derives from spatial scales, one might imagine that it is attributable to the way the samples were selected, as spatial blocks of particular sizes, and that it might disappear if samples of settlements were randomly selected individually from the total population of settlements, satisfying much more fully the fundamental assumptions of random sampling. A Monte Carlo experiment to pursue this possibility selected 1000 samples of 30, 60, 120, and 240 settlements from the same population. These sample sizes were chosen to approximate the numbers of settlements included in the spatial blocks of different sizes in the previous experiment. The results of this experiment are presented in Table 2 for comparison to those of the first experiment.

Table 2
Means and standard deviations for A values for samples of different numbers of settlements from the randomly generated population of 1452 settlements

Sample size	Number of samples	Mean A	Standard deviation
30 settlements	1000	0.636	0.071
60 settlements	1000	0.706	0.046
120 settlements	1000	0.718	0.046
240 settlements	1000	0.765	0.027

While the average A values are higher, and thus closer to the "true" A value for the entire population, they do not provide good estimates. Even the highest average A value (from the 1000 240-settlement samples) has a standard deviation that does not include the population value of 0.796. Multimodality in the distributions of A values for this second Monte Carlo experiment is less pronounced than in the first, but definitely present, and owing to the same underlying cause. In short, this experiment is as discouraging about rank-size graphs produced from samples as was the first.

These two Monte Carlo experiments, of course, were based on an artificial population without the kind of hierarchical structure the rank-size rule would encourage us to expect (hence the markedly convex rank-size pattern in Fig. 4 and the A value of 0.796). For further experimentation, two levels of settlement hierarchy were added to this population by superimposing a perfect hexagonal lattice of 11 top-rank "central places" and 42 second-level centers (Fig. 6). The top-rank centers range from 22 to 81 ha, with a mean of 44.5 ha, and the second-level centers range from 6.2 to 46 ha, with a mean of 18.9 ha. There is thus a clear difference in settlement size between the tiers, although there is substantial overlap, as one might expect in a real dataset. (It is worth noting in passing that the common practice of identifying settlement hierarchy by seeking multimodality in a settlement-size histogram without regard to settlement location would produce wholly

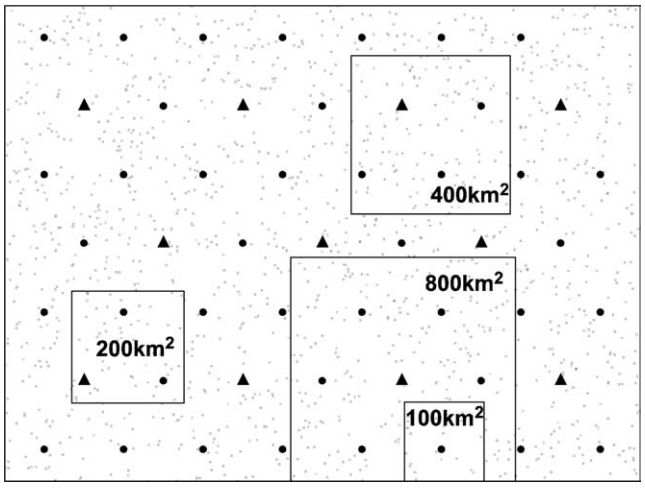

Fig. 6. Hexagonal lattice of two levels of centers superimposed on randomly generated population of settlements, with sample blocks of the four different sizes employed in the Monte Carlo experiment. (Triangles indicate highest-level centers; circles, second level settlements.)

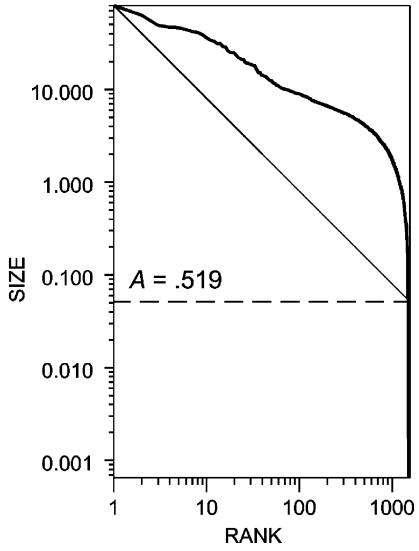

Fig. 7. Rank-size graph for artificial population of settlements with hexagonal lattice of centers.

inappropriate results in any dataset characterized by such an overlap in sizes between settlements of different tiers.) This artificial population produces a rank-size pattern that, while still convex, is closer to log-normal than that of the purely random population (Fig. 7). This closer approach to log-normality is reflected in an A value closer to zero: 0.519 (compared to 0.796 for the random population). The results of a Monte Carlo sampling from this more structured population of settlements, carried out exactly as in the first experiment described above, are shown in Table 3. Average A values for all sample block sizes depart even farther than before from the A value for the entire population. Once again,

Table 3
Means and standard deviations for A values for sample blocks of different sizes from the randomly generated population of 1452 settlements with the addition of a hexagonal lattice of higher order centers

Sample block size (km^2)	Number of samples	Mean A	Standard deviation
100	1000	0.111	0.306
200	1000	0.214	0.214
400	1000	0.271	0.118
800	1000	0.379	0.075

none of the sample block sizes would provide a useful estimate of the population value. Consideration of just how sample blocks of different sizes can fall on to the structure of the hexagonal lattice makes even clearer the reasons behind the peculiar behavior we have observed. For example, a 100 km^2 block (Fig. 6) will always include an assortment of lower-tier settlements along with one higher-level center (or very rarely, two). Consequently such a sample block will usually have one rather large settlement and show either a primate or a convex pattern, depending on just how much smaller the second largest settlement is. A values will range from strongly negative to strongly positive, averaging out to the near log-normal value of 0.111 for 1000 samples of this block size. In addition to an assortment of lower-tier settlements, a 200 km^2 block will contain two, or more commonly, three higher-order settlements. Four or five such centers will fall within the boundaries of a 400 km^2 sample block, and a block 800 km^2 in size will capture six or eight higher-order settlements along with a larger number of lower-tier settlements. This increasing number of higher-order settlements captured by larger block sizes will result in increasingly convex rank-size patterns (and increasing positive average A values) as block size goes up. With this hexagonal lattice of higher-order centers it becomes quite obvious how samples of different sizes characterize qualitatively different patterns that occur at different spatial scales. The same phenomenon occurs regardless of whether the samples consist of all the settlements in a spatial block or of settlements individually selected at random. The size of the sample inevitably determines the average nature of the set of settlements included and this corresponds to the spatial scale of the analysis.

Thinking about sampling in this way is just not applicable to rank-size analysis, where defining the boundary of a region determines the phenomenon under study. As an example, the entire area illustrated in Fig. 6 bounds what we might imagine as a settlement system of 11 separate units, each consisting of a major center, several second-level centers, and a larger number of third-level settlements. Analyzed at this scale there is little overall integration, and this is reflected in the

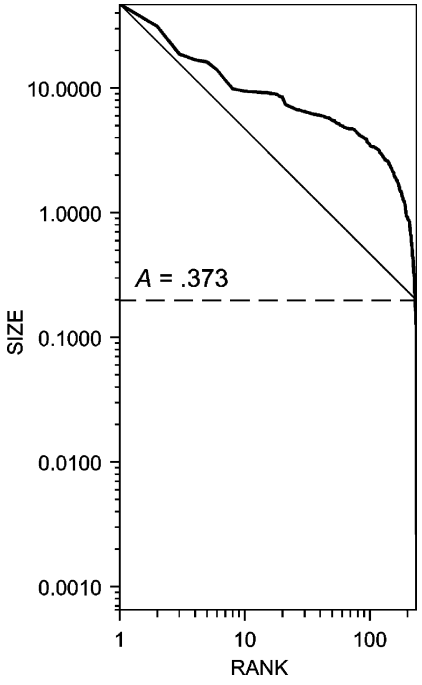

Fig. 8. Rank-size graph for the settlement system delimited by the 800 km^2 block in Fig. 6.

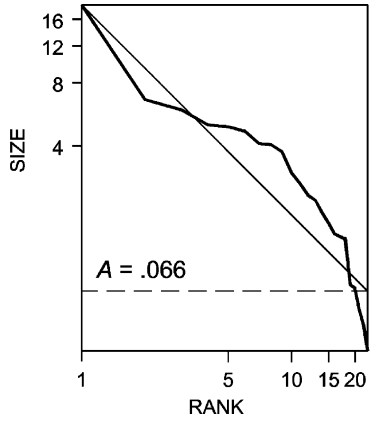

Fig. 9. Rank-size graph for the settlement system delimited by the 100 km^2 block in Fig. 6.

convex rank-size pattern in Fig. 7, and the *A* value of 0.519. Johnson [9 (pp. 498–499), 10 (pp. 240–242), 11 (pp. 167–176)] and Paynter [19 (pp. 253–260)], who have paid particular attention to this issue, would identify this as "pooling" of multiple settlement systems. At a smaller scale, the 800 km^2 block shown in Fig. 6 roughly delineates one of the 11 separate units, producing a case where, in Paynter's [19 (p. 254)] words, "analytical and theoretical boundaries coincide." We would likely imagine this unit to show a fairly high degree of internal integration, and this is reflected in its rank-size pattern (Fig. 8) and its *A* value of 0.373. The next two smaller-sized sample blocks (400 km^2 and 200 km^2) would not contain a theoretically meaningful settlement unit, no matter where they were placed, because of what Johnson [10 (pp. 240–242)] would call "partitioning" of a larger system. At the smallest scale, a 100 km^2 block might again contain a meaningful settlement unit: one center (either first or second level) and a set of subsidiary settlements. Analysis of the 100 km^2 block shown in Fig. 6 produces the rank-size graph in Fig. 9 and an *A* value of 0.066—the smallest departure yet from the log-normal pattern. These are simply different analyses of spatial patterning at different scales; it makes no sense to try to think of the smaller blocks as samples useful for characterizing the larger population.

4. Statistical significance: the bootstrap

The recognition that, for rank-size analysis, a survey tract cannot meaningfully be considered a sample from a larger region might be taken to imply that the issue of statistical significance does not arise. This implication, while superficially seductive, is incorrect, for reasons explained cogently by Cowgill [5 (pp. 366–367)]. In a nutshell, larger numbers of observations provide more convincing support for conclusions than smaller numbers of observations do, and larger differences between sets of observations provide more convincing evidence of different underlying dynamics. Just how large the numbers of observations must be and how large the differences between sets of observations must be to justify a particular level of confidence is the heart of statistical inference, whether the sets of observations are literally samples or not. We may, for example, compare the rank-size patterns of two settlement systems and conclude that different settlement dynamics are represented. In such a case, the greater the number of settlements in each system and the larger the difference between them, the greater our confidence in our conclusion should be. Making useful comparisons requires a means of saying how large a difference and how many observations are required for what level of confidence. Previous attempts to accomplish this in rank-size analysis have taken the Kolmogorov–Smirnov test and comparison to the log-normal pattern as a basis [18,21]. The *A* coefficient proposed above offers a different, more versatile, and more satisfying approach.

Standard parametric statistical approaches are of no use for evaluating the significance of differences in *A* values between two sets of observations because we have no theoretical basis for knowledge of the sampling distribution of *A*. This is precisely the situation for which resampling approaches, such as the bootstrap and the jackknife, were proposed. These techniques have demonstrated their effectiveness in situations that are well known theoretically, as in estimating means, and have been applied in a wide variety of contexts for which, like rank-size analysis, there is no theoretical basis for characterizing sampling distributions [17].

These methods have been used in archaeology, for example, for giving diversity indexes error ranges for different confidence levels [1,12]. Resampling approaches are especially appealing for rank-size analysis because they require no theoretical assumptions about the population samples are drawn from—indeed they do not require imagining such a population at all. The bootstrap produces an approximation of the sampling distribution by repeatedly selecting samples from the initial set of observations. These samples have the same number of observations as the original set, but vary because the sampling is done with replacement. Thus each sample may duplicate some observations from the original set while omitting others. The bootstrap is used here instead of the jackknife because the bootstrap utilizes a large number of samples even when the number of observations is small.

4.1. The bootstrap and early Taosi settlement (North China)

The application of the bootstrap to the A coefficient can perhaps best be illustrated with rank-size analysis of actual settlement datasets. First we will look at the late Neolithic Taosi (2500–2000 BC) dataset from the Linfen Basin in Shanxi Province, north China as presented by Liu [15,16]. A rank-size plot for 29 early Taosi settlements can be seen in Fig. 10. The A value of -0.512 reflects the fact that many settlements are smaller than the log-normal expectation. At the same time, we would not characterize this as a primate pattern because the distribution is not dominated by a single large settlement, but rather, the several largest settlements adhere roughly to the log-normal line. We might well wonder how confident we should be in reaching the conclusion that this dataset indicates an underlying settlement dynamic at all different from the "expected" log-normal pattern.

To assess our confidence with the bootstrap, we randomly select (with replacement) 1000 samples of 29 from the original dataset of 29 settlements. For each of the 1000 samples, the value of A is calculated. These 1000 A values are represented in the histogram shown in Fig. 11. As the Monte Carlo experiments would lead us to expect, this distribution is asymmetrical, displays a tendency toward multimodality, and has a mean value which differs sharply from the value of A calculated for the original dataset (-0.294 compared to -0.512). It is common to find such a difference between the A value calculated for a set of settlement sizes and the mean A value for the 1000 bootstrap samples selected from it. On average, the bootstrap samples produce rank-size curves closer to the log-normal line than does the original dataset, so the mean A value for the bootstrap samples is generally closer to zero than is the A value in the original dataset. The rank-size pattern is, of course,

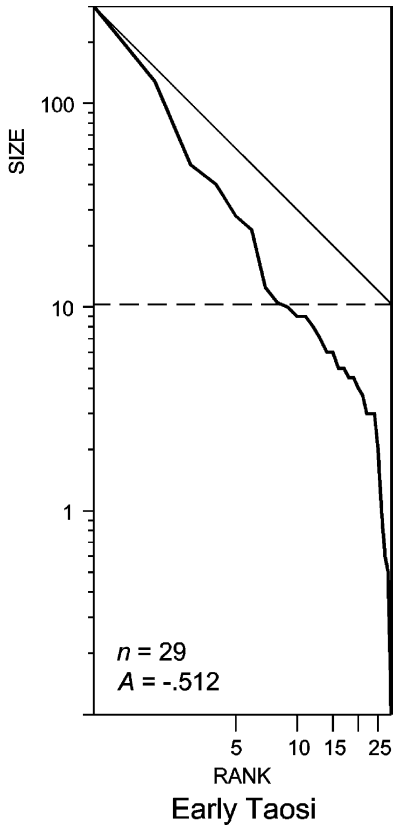

Fig. 10. Rank-size plot of early Taosi phase settlements.

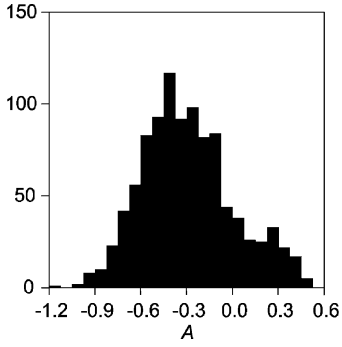

Fig. 11. Histogram of 1000 A values from bootstrapping the early Taosi dataset.

characterized by the A value calculated directly from the original dataset. The distribution of A values from the bootstrap samples is only used to approximate the spread of the sampling distribution of A. If this distribution of 1000 A values had a normal shape we could readily calculate its standard deviation and use this as the basis of assigning an error range to A at the desired confidence level. The standard deviation, however, would not be a useful measure for such an asymmetrical, and even multimodal, distribution. This is commonly the case in situations where the bootstrap is applied, and

a number of approaches to establishing error ranges have been suggested. The percentile method [17] seems most applicable to rank-size analysis. The lower and upper cut points for a 90% confidence error range, for instance, are the 5th and 95th percentiles of the distribution of A values from the 1000 bootstrap samples; for a 95% confidence error range, the 2.5th and 97.5th percentiles; and so on.

In order to utilize the percentile method in rank-size analysis where the mean A from the 1000 bootstrap samples differs from the A value from the original dataset, the A distribution from the bootstrap samples must be transformed. A transformation based on the t distribution has been suggested for this purpose [17], but it requires assumptions of normality that seem decidedly out of place in this context. A simpler arithmetic transformation seems more appropriate for rank-size analysis. The A value from the original dataset is subtracted from the mean A value for the bootstrap samples; this difference is added to each of the 1000 A values for the bootstrap samples before the percentile cut points are determined. It is as if the histogram in Fig. 11 were simply shifted to the right so that its mean would be the same as the A value from the original dataset. Applying this technique to the early Taosi dataset, we finally state its A value as between -0.957 and 0.066 with 90% confidence. Since the A value for the log-normal line is 0.000, which falls barely within this 90% confidence level error range, we are slightly less than 90% confident that the early Taosi data represents a settlement dynamic different from that suggested for log-normal. This is, of course, equivalent to saying that there is only moderate significance to the difference between the observed early Taosi settlement data and the expected log-normal pattern; the significance probability P would be somewhat greater than 0.100.

The A coefficient expressed with an error range for the desired confidence level, then, helps us characterize rank-size patterns for comparison to log-normal and to evaluate the significance of any difference observed. Finally, we can connect this numerical analysis back to the intuitive simplicity of rank-size graphs. The graph in Fig. 12 adds an approximate 90% confidence zone for the rank-size curve to the original presentation in Fig. 10. This zone is delimited by two additional rank-size curves produced by plotting the settlement sizes for two of the 1000 bootstrap samples. The upper line represents the sample whose A value was closest to the upper cut point of the 90% confidence error range; the lower line, the sample whose A value was closest to the lower cut point. Both samples were required to satisfy an additional condition: the largest settlement had to be included so that all the lines on the graph would originate from the same point. This zone is very much an approximation, but it gives a sense of how precisely we can characterize the underlying settlement dynamic on

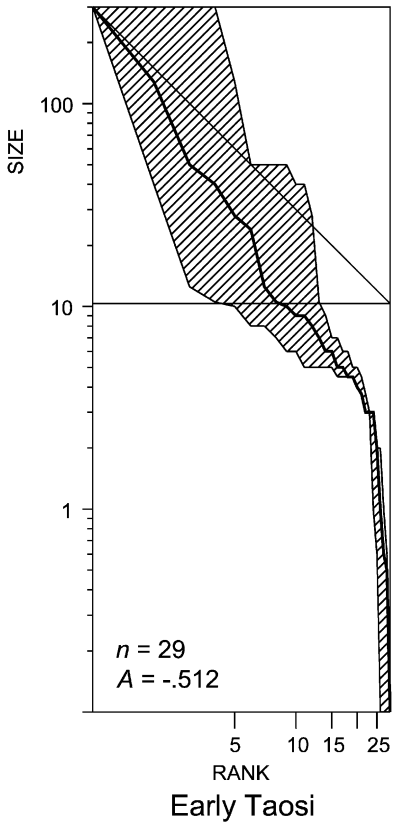

Fig. 12. Rank-size plot of early Taosi phase settlements with approximate 90% confidence zone for rank-size curve.

the basis of the observations we have. If the zone adheres tightly to the rank-size curve, then we have high confidence that the rank-size pattern observed paints a fairly good picture of this dynamic. If, on the other hand, the zone only loosely conforms to the rank-size curve, then we must recognize that our observations permit only a fuzzy view of settlement dynamics.

A program to calculate A with bootstrapped error ranges is available from the authors <www.pitt.edu/~drennan/ranksize.html>. It is written in FORTRAN for Windows operating systems and is designed to be used in conjunction with the graphing tools of a general-purpose statistical software package (e.g. SYSTAT, SPSS, etc.).

4.2. The bootstrap versus the Kolmogorov–Smirnov test

Using an approach suggested by Savage [21], Liu [16] determined that the early Taosi dataset represents a highly significant departure from log-normality, with a P value less than 0.001, compared to the P value slightly greater than 0.100 from our analysis. This difference arises because the logic of Savage's approach to significance in rank-size analysis is fundamentally different from ours. Its frame of reference is Monte Carlo sampling, which, as we have seen, is confounded by the

Table 4
A values and bootstrapped error ranges for perfect log-normal settlement systems with different numbers of observations

Number of settlements	*A*	Error range (90% confidence)
8	0.000	0.925 (−0.470 to 0.455)
15	0.000	0.774 (−0.374 to 0.400)
30	0.000	0.678 (−0.318 to 0.360)
60	0.000	0.583 (−0.253 to 0.330)
120	0.000	0.488 (−0.205 to 0.283)
240	0.000	0.432 (−0.171 to 0.261)

implications of spatial scales peculiar to rank-size analysis. Moreover, Savage's approach depends almost entirely on the sample fraction, conceived as the proportion of ancient settlements actually discovered by archaeological survey. As Savage [22] notes, however, this approach almost inevitably produces high significance values for sample fractions greater than 80% and low values for sample fractions less than 70%, regardless of the number of sites in the sample. But, as we have seen above, statistical confidence in rank-size analysis is not based on literal relationships between samples and populations but rather relates to our confidence in broader interpretations of the observations we have made. The imaginary populations involved in this logic can only be conceived of as limitless. Obviously, the notion of sample fractions is meaningless for such statistically infinite populations [4 (p. 263)]. As is generally recognized in statistics, then, confidence must depend primarily on the number of observations in the sample rather than the proportion of the population in the sample. In the present context, this means that we should have little statistical confidence in discussing rank-size patterns based on very few settlements, even if we could be certain that they constituted 100% of the settlements in a study region. The *A* coefficient proposed here gives us bootstrapped error ranges that narrow in as the number of observations increases (for a given confidence level), just as common sense and statistical principles dictate (see example in Table 4). The reverse implication of the error ranges in Table 4 is that a rank-size pattern based on eight observed settlements would have to differ from log-normal by more than twice as much as one based on 240 observations for us to have 90% confidence that a departure from log-normality was really indicated.

We agree with Savage [21 (p. 243), 22 (p. 11)] that completeness of survey coverage is a vital consideration in rank-size analysis and that total recovery is virtually never achieved so that we are always working with samples of settlements. If a rank-size graph is produced for substantially less than 100% of the settlements in a region, though, the result is not just lacking in significance. Instead, the graph characterizes, not the entire region at all, but a completely different smaller-scale socio-spatial phenomenon, as we saw in the Monte Carlo experiments. This is a substantive issue, not to be dealt with after the fact with a significance test, but to be resolved in the initial definition of the object of study. Statistical confidence more properly concerns the process of making broader conclusions about the nature of this object as observed in regional settlement data. As such, it depends, not on the number of sites missed on survey, but on the number of settlements observed.

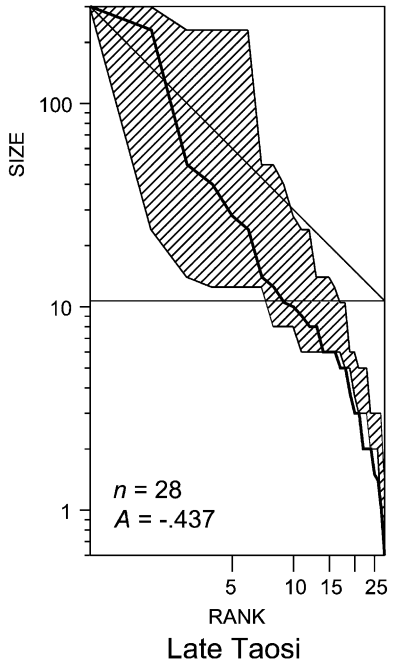

Fig. 13. Rank-size plot of late Taosi phase settlements with approximate 90% confidence zone for rank-size curve.

4.3. The bootstrap and comparison: early and late Taosi settlement systems

The Taosi dataset further enables us to extend our exploration beyond the synchronic characterization of a single settlement system to the study of change through time. Here the comparison of observed rank-size patterns to log-normal is less useful than direct comparison of the patterns from different periods to each other. The late Taosi settlement rank-size graph with a 90% confidence zone is shown in Fig. 13 for comparison to Fig. 12. This pattern yields an *A* coefficient of −0.437. This time, we have to lower the confidence level below 80% before the error range excludes the 0.000 value for the log-normal line. We are thus somewhat less than 80% confident that the late Taosi settlement system departs from log-normality. More importantly, the difference between early and late Taosi *A* values is so small (0.075) in relation to error ranges for various confidence levels (Fig. 14) that, irrespective of any comparison to

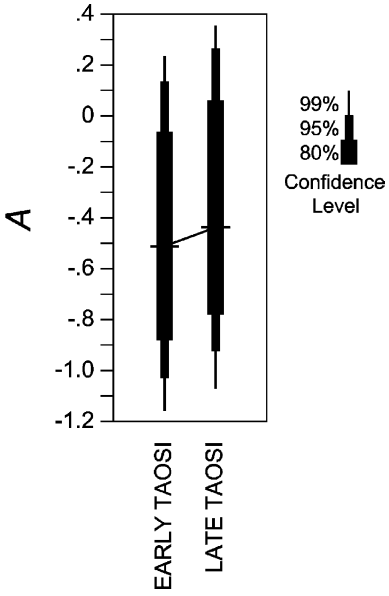

Fig. 14. *A* values for early and late Taosi phase settlement systems compared.

log-normal, we would have virtually no confidence in concluding that any change had occurred in this aspect of settlement organization. Again these are conclusions that differ sharply from those reached by Liu [16 (pp. 247, 250–251)], following a different approach to evaluating the significance of rank-size patterns. The *A* values of −0.512 and −0.437 unambiguously identify rank-size curves whose shape departs from log-normal in the same direction for both periods, in contrast to Liu's subjective identification of early Taosi as a primate pattern and late Taosi as convex. The error ranges for *A* and the confidence zones in Figs. 12 and 13 emphasize how similar the two rank-size patterns are overall and how little significance the difference between them has. The analytical approach proposed here, then, suggests that two sets of 29 and 28 settlements, respectively, with rank-size curves no more different than these are, do not give us enough confidence to identify change in settlement organization between early and late Taosi. This might, of course, have worked out differently. If the rank-size curves had differed more strongly in shape, then the difference between the *A* coefficients would have been greater. If this direct indication of a difference in shape had been strong enough, it would have turned out to have some statistical significance, even though the number of observations was small. On the other hand, if early and late Taosi had been represented by larger numbers of observed settlements, even the weak difference in rank-size shapes we observed could have achieved a high degree of significance. As with other significance tests, large samples can lead to the conclusion that a difference so weak as to have little conceivable interpretive value is none the less of statistical significance. Varied datasets from several other parts of the world further illustrate the versatility of these tools.

5. Applications of *A* with error ranges

5.1. Susiana Plain (Southwestern Iran)

In a pioneering application in archaeology, Johnson [8–11] used rank-size graphs to delineate increasing centralization in the Susiana Plain of southwestern Iran in Uruk times (3900–3200 BC). The sequence of rank-size graphs for the four periods (Fig. 15) is summarized by the decreasing *A* values with error ranges graphed in Fig. 16. The error ranges show that, in general, when we compare any period's rank-size pattern to that of the immediately preceding or following period, we would have about 70% confidence in concluding that a change in underlying settlement dynamic had taken place. Clearly, however, the overall trend throughout the period, just as Johnson recognized, is one of decreasing convexity. When we look at the cumulative change over the entire sequence of four periods, we have greater than 95% confidence in discussing a fundamental shift in underlying settlement dynamic from Terminal Susa A to late Uruk. That is, there would be less than a 5% chance that we had been merely led astray by an insufficient number of observations. The settlement data presented by Johnson make it seem that this change was relatively gradual throughout the period involved, but we would have only about 70% confidence in this more detailed conclusion. The exact pacing of the overall change could be charted with greater precision and confidence if there were more observations for each period. In this instance, use of the *A* coefficient with error ranges has not led to any different conclusion than originally proposed, but it has given us a clearer sense of how confident we should be in predicating discussion of sociocultural change on these results.

5.2. Santa Valley (Northern Peruvian Coast)

In his settlement study of the Santa Valley, Peru, Wilson [23,24] did not use rank-size analysis, but his careful attention to demographic reconstruction provides an uncommonly sound foundation for it: settlement size here is not just archaeological site area but an actual population approximation. The rank-size graphs in Fig. 17, summarized in Fig. 18, detail changing patterns from first sedentary living about 2100 BC until the Moche conquest around AD 400. The small number of observations for the preceramic Las Salinas occupation makes its rank-size pattern difficult to characterize with much confidence at all, as shown by the very large error ranges in Fig. 18. Following an occupational hiatus, full-blown agricultural occupation in Cayhuamarca shows a distinctly convex rank-size pattern, supporting, and lending substantial confidence to, Wilson's

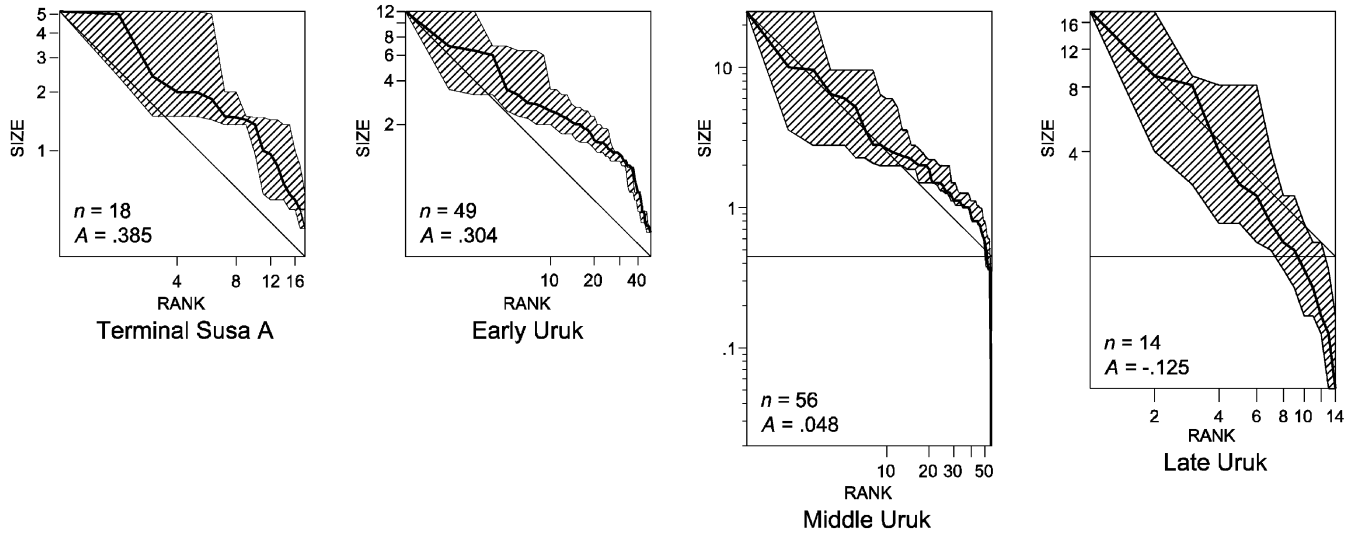

Fig. 15. Rank-size graphs with 90% confidence zones for Uruk period settlement on the Susiana Plain.

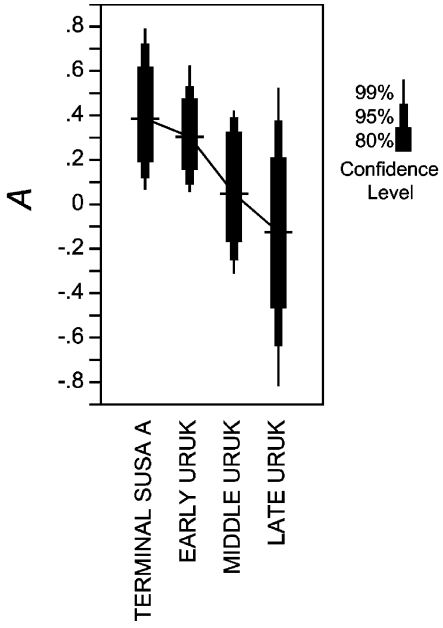

Fig. 16. *A* values for Susiana rank-size patterns.

[24] reconstruction of valley settlement as several largely autonomous clusters. The Vinzos rank-size pattern tends more toward primateness, differing strongly from that of Cayhuamarca, and the shortness of even the 99% confidence level error bars depicted in Fig. 18 (compared to the distance between the *A* values for the two phases) tells us that this difference has considerable statistical significance [7 (pp. 151–153)]. There is a settlement system change here that merits discussion; it is not, however, simply a case of increased regional integration, as more primate rank-size patterns are often taken to indicate. The "primate" settlement that produces this pattern is not an emerging civic-ceremonial center, but instead a village at the southwestern extreme of the valley, removed from other settlement clusters, and undistinguished from other habitation sites except by its unusually large size. The rank-size analysis for this phase, then, calls attention to a shift in settlement dynamics which Wilson's dataset makes it possible to discuss with confidence, even if it is not interpretable in the familiar terms of increasing integration. The Vinzos phase underscores the risk of over-interpretation of rank-size graphs taken alone. They become most meaningful with reference to which specific sites are at the top of the graph, where they are located in the region, and what their other characteristics are [11 (p. 177)].

Continuing with the Santa Valley sequence, there is little statistical significance to the change in rank-size pattern from Vinzos to early Suchimancillo (Fig. 18), but the near-zero *A* value for early Suchimancillo, unlike that for Vinzos, does appear to reflect regional integration, in just the way such log-normal patterns are usually interpreted. Wilson [23 (p. 237)] identifies fewer settlement clusters than before, "local centers" appear for the first time, and the two largest settlements at the top of the rank-size graph are local centers in the two principal settlement clusters which encompass almost the entire population of the valley. Although there is not much significance to the difference in *A* values between Vinzos and early Suchimancillo, Fig. 18 shows that we can have very high confidence that a settlement dynamic different from that of the Cayhuamarca phase had emerged by early Suchimancillo. The rank-size pattern for late Suchimancillo is again distinctly convex, and Fig. 18 gives us a moderate degree of confidence in postulating a change in settlement dynamic from early Suchimancillo. Most of the largest settlements at the top of the rank-size

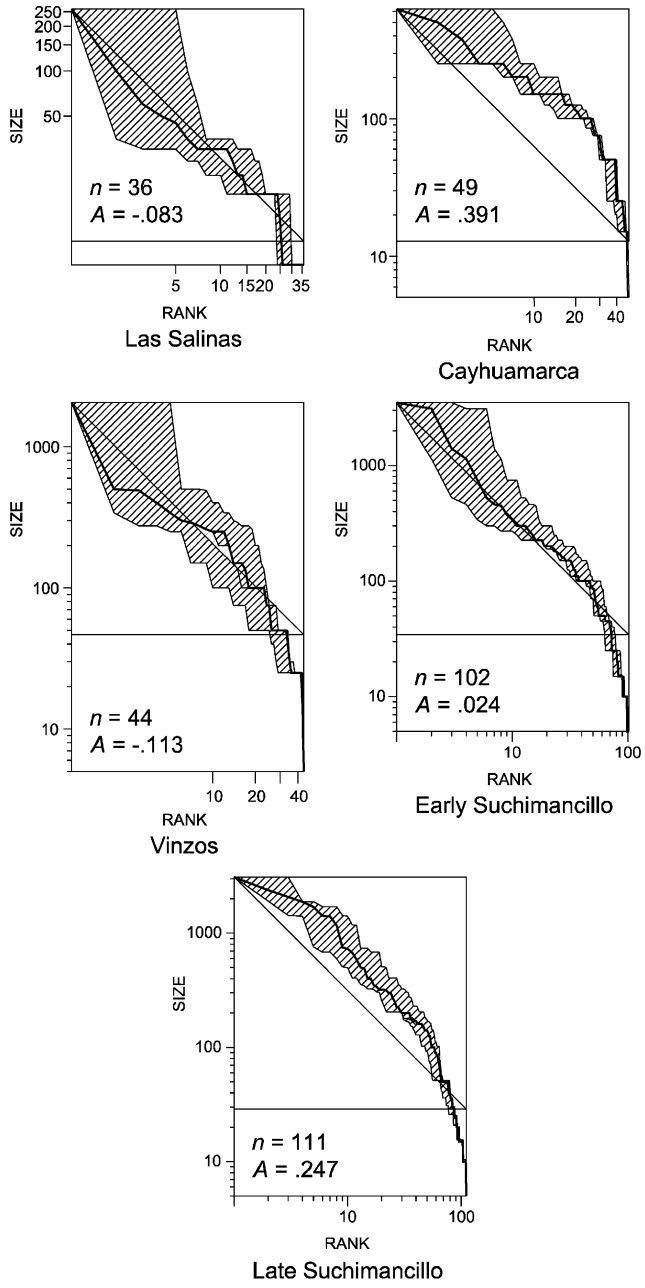

Fig. 17. Rank-size graphs with 90% confidence zones for the Santa Valley sequence.

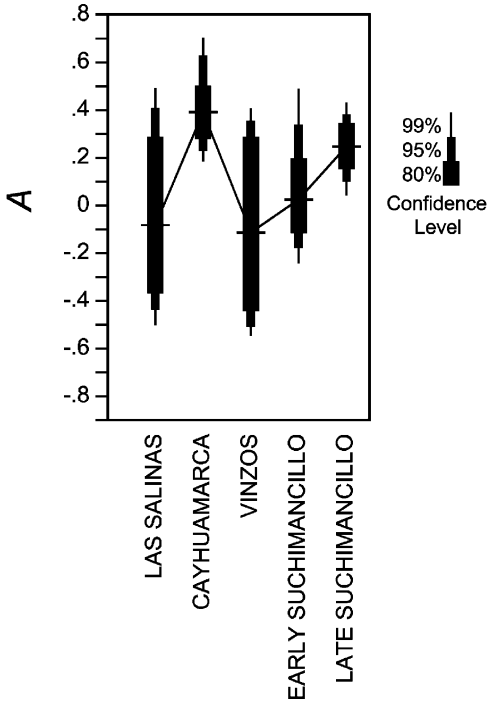

Fig. 18. *A* values for Santa Valley rank-size patterns.

graph are identified by Wilson [23 (p. 241)] as local centers, and the convex pattern reflects the increase in their numbers over early Suchimancillo times. No regional center has emerged that really dominates the centralized local units that emerged in early Suchimancillo times; rather, an increased number of such units is packed into the same region, in what Wilson [24 (p. 355)] describes as a "loosely integrated chiefdom society." Rank-size analysis might be applied separately to each local unit to delineate and characterize these smaller-scale phenomena [11 (p. 169)]. As discussed above, this would refocus the analysis on settlement patterning at a smaller scale, and would not comprise "sampling" from the larger region. Our analysis ends with late Suchimancillo because rank-size analysis for subsequent periods is confounded by a lack of correspondence between archaeological sites and human settlements. Aggregation of archaeological sites into larger units would be required for meaningful analysis, and this is beyond the scope of the present paper.

5.3. Valley of Oaxaca (Southern Mexican Highlands)

Blanton et al. [3] and Kowalewski et al. [14] performed precisely such aggregation on the settlement data from the Valley of Oaxaca, Mexico, so as to delineate units more appropriate for rank-size analysis (based on actual population estimates for each settlement). Figs. 19 and 20 provide the essence of a rank-size analysis of this sequence from the establishment of sedentary living around 1500 BC to the Spanish Conquest at AD 1520. The earliest settlement system in the Tierras Largas phase shows a strongly convex rank-size pattern, which shifts to a markedly primate pattern in the San José phase, focused on the burgeoning center of San José Mogote. Not surprisingly, this very strong change also has extremely high statistical significance (Fig. 20). Somewhat more surprisingly, perhaps, the lengths of the error bars in Fig. 20 for the two phases differ dramatically, even though the number of observations is relatively small for each. We might well expect

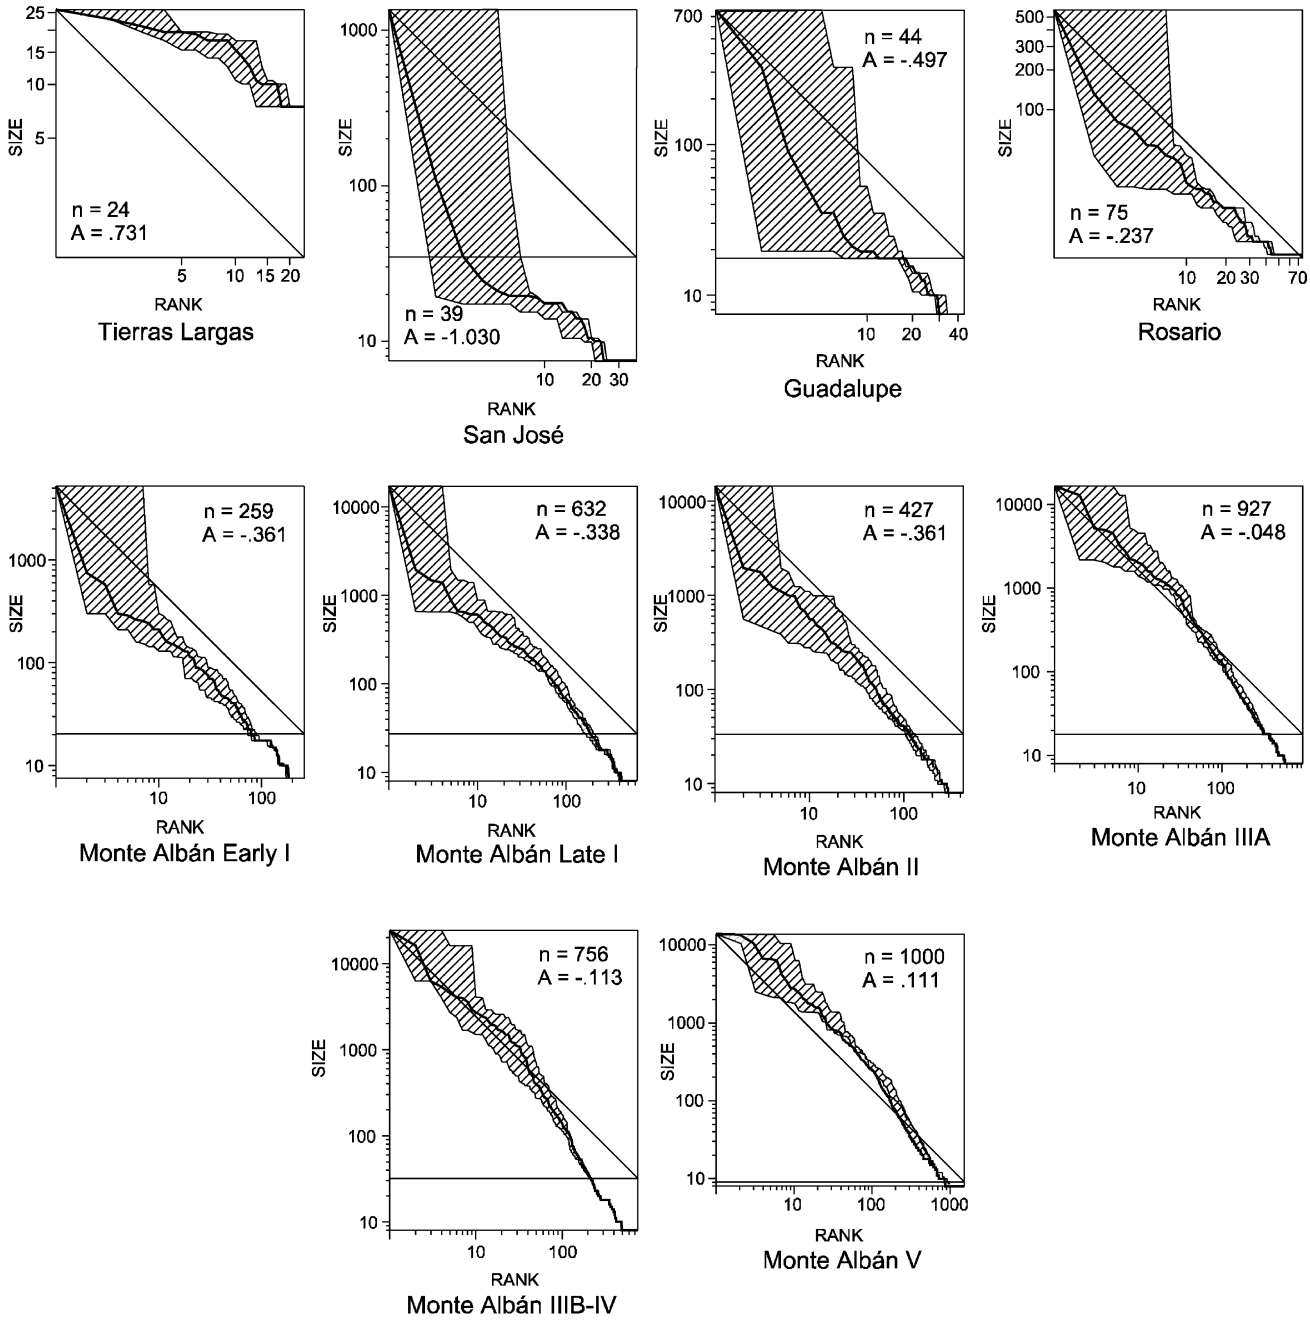

Fig. 19. Rank-size graphs with 90% confidence zones for the Valley of Oaxaca sequence.

such small numbers of observations to produce fairly large error ranges—larger for Tierras Largas with only 23 settlements than for San José with 39. The comparison between these two phases, however, illustrates the impact that the shape of the rank-size curve has on statistical confidence. Strongly primate patterns, like that of San José, regularly produce large error ranges, even for quite large numbers of observed settlements. While it might seem that the bootstrap is malfunctioning in this context, it actually makes very good sense. This is because, regardless of the total number of settlements observed, primateness is produced by a very few settlements in the topmost ranks, often depending almost entirely on the size difference between the first and second ranked settlements. Such dependence on a very few observations is appropriately reflected by the large error ranges the bootstrap produces for such shapes. A more convex pattern, like that of Tierras Largas, is not contingent on such a tiny number of observations in the top ranks, so error ranges for convex patterns are usually smaller, even when the total number of observations is not large. Assessing statistical confidence with

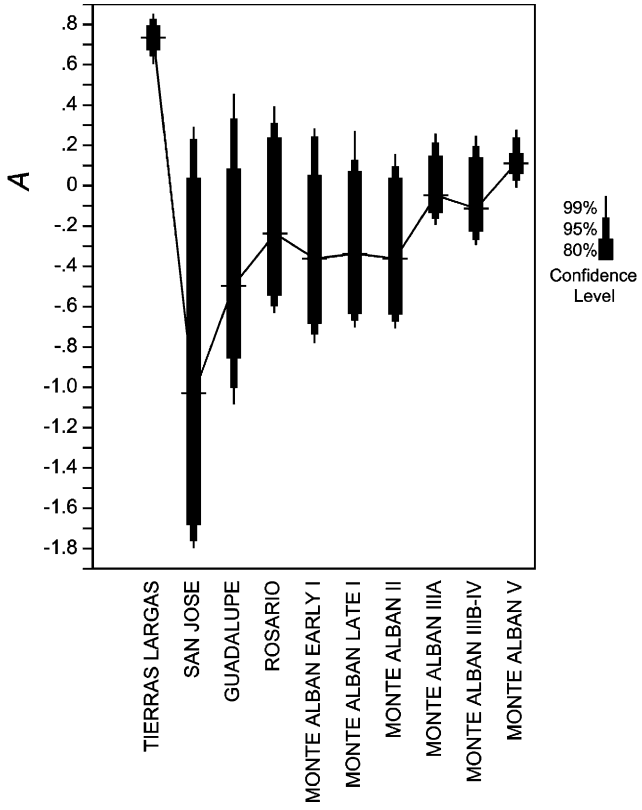

Fig. 20. *A* values for Valley of Oaxaca rank-size patterns.

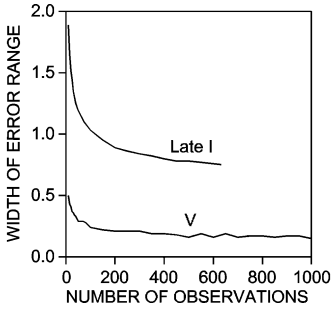

Fig. 21. Size of 90% confidence error ranges for Monte Albán Late I and V in relation to number of observations.

bootstrapped error ranges for the *A* coefficient, then, appropriately handles the subtleties of the rank-size pattern's dependence on a few high-ranked settlements within the dataset, as this dependence varies according to the shape of the curve.

The next five phases continue through the founding and development of the regional capital at Monte Albán, and are also characterized by primate rank-size patterns, although not as strongly as the San José phase. Fig. 20 tells us that we can have considerable confidence in describing a relaxation over time of the forces that produced an intensely primate pattern in San José times. These less strongly primate patterns are also the product of larger numbers of observations; consequently the error ranges are smaller. Nevertheless, the error bars for even low confidence levels are long compared to the differences between *A* values for these phases, so we would have little confidence in identifying changes in basic settlement dynamic from Rosario through Monte Albán II [cf. 14 (pp. 77, 100–101, 146, 187)]. Guadalupe occupies a transitional position between San José and the later, significantly less primate patterns. The next major change in rank-size patterns is between Monte Albán II and IIIA. As Kowalewski et al. [14 (p. 230)] note, the IIIA pattern is substantially more log-normal than for any previous period, and Fig. 20 shows that we should indeed have high confidence in identifying a change at this point in the sequence. Despite the decline of Monte Albán as regional center after IIIA, Fig. 20 shows that the rank-size pattern for the valley does not change significantly [14 (p. 290)]. Although Jalieza, the IIIB-IV top-rank settlement, is a very different kind of community with far less public architecture than IIIA Monte Albán, its position in the rank-size hierarchy of its period is quite similar. It is at the beginning of Monte Albán V that we can again confidently identify a shift in settlement dynamic, reflected in a significantly more convex rank-size pattern (Fig. 20) interpreted as indicating less regional integration than at any time since the Tierras Largas phase [14 (p. 324)].

The wide range in numbers of settlements through the periods in the Oaxaca sequence enables us to explore systematically the impact of number of observations on bootstrapped error ranges. Error ranges were calculated repeatedly for the Monte Albán Late I and V *A* values, for decreasing numbers of settlements by progressively omitting larger and larger groups of the lowest-ranked settlements from each set. Fig. 21 shows how the error range for each period decreases as the total number of observations increases. Illustrating the effect of curve shape on error ranges as discussed above, the error ranges for the very primate Monte Albán Late I curve are substantially larger than those for the convex Monte Albán V pattern for any number of observations. As always in evaluation of confidence, when the number of observations is small, even modest increases in the number of observations narrow the error ranges appreciably. For Monte Albán Late I, the 90% confidence level error range decreases as the number of observations goes up, sharply at first, then more gradually until about 400 observations is reached, beyond which there is little gain in confidence up to the full 632 settlements in the dataset. For the more convex Monte Albán V pattern, the error ranges decrease only slightly beyond about 150 observations, and increasing the number of observations above 500 makes virtually no difference in confidence. Although the entire dataset includes over 1500 settlements, the bootstrap was applied only to the 1000 largest, since computing time

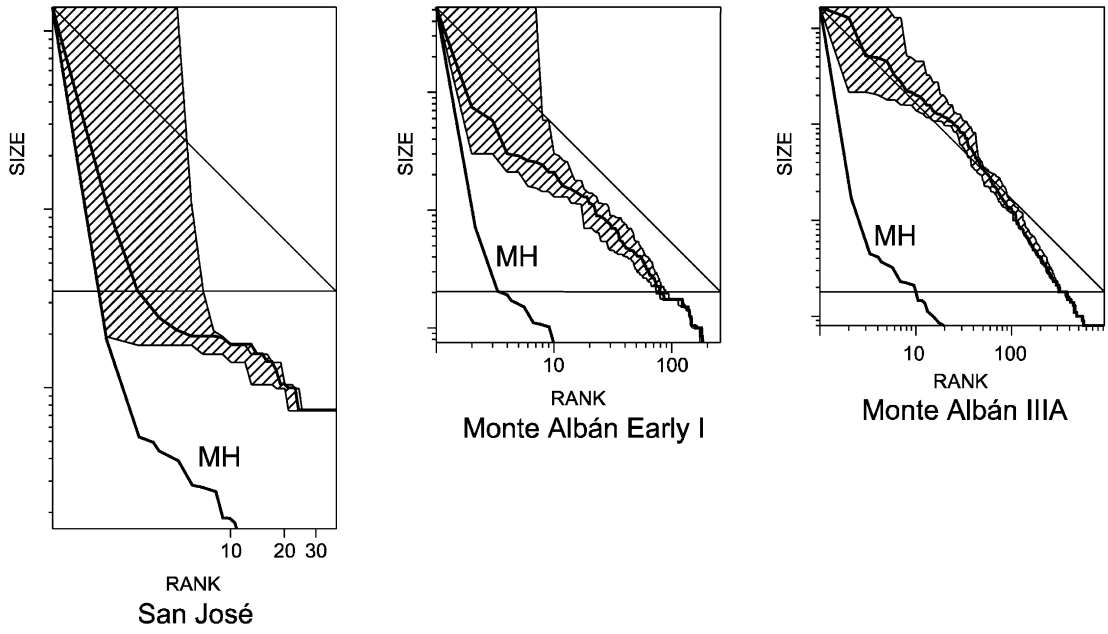

Fig. 22. Rank-size curve for Middle Horizon (MH) Basin of Mexico [2 (p. 196)], overlaid on rank-size graphs with 90% confidence zones (shaded) for three periods in the Valley of Oaxaca.

for resampling increases sharply beyond this point, and it is clear that the results would not be meaningfully different. It seems unlikely that there would ever be need to include more than 1000 observations in any analysis, and this recalls the practice sometimes followed of performing rank-size analysis on only the higher ranking settlements. Inclusion of as few as the largest 20 to 30 settlements often does produce an accurate characterization of the shape of the curve, since it is the largest settlements that matter most. Clearly, however, substantially more are needed for a good assessment of statistical confidence—likely the entire dataset unless it is extremely large.

6. Conclusion

In these example applications of the *A* coefficient with bootstrapped error ranges, we have focused on the comparison of rank-size patterns for different periods within a single region to each other, downplaying deviation from the log-normal "expectation" that is often the center of rank-size analysis. The *A* coefficient enables us to make these comparisons more precise, although this comes at the cost of a potentially dangerous degree of abstraction. *A* is useful, not as a substitute for visual examination of rank-size graphs, but as a complement to such examination, just as rank-size graphs themselves are most effectively interpreted in conjunction with other kinds of information about settlements, their characteristics, and their distributions [11 (p. 177)]. Combined with the bootstrap, *A* offers a meaningful way to assess the confidence we should have in the rank-size patterns we observe, thereby concentrating our interpretive energy on changes in settlement systems that we are confident actually occurred—instead of finding reasons for differences that may represent nothing more than the vagaries of sampling. The rank-size analyses of settlement data presented here focus on change through time, but *A* with error ranges can just as well be applied to inter-regional comparisons. Such comparisons are likely to rely at least in part on published information, but unfortunately some authors present only their rank-size graphs and not the data needed to recreate them, calculate *A* values, or determine error ranges. Rank-size plots with graphical confidence zones, however, can be compared to summary plots for other regions, even when full supporting data are not available. Fig. 22 shows that Teotihuacan-dominated Middle Horizon Basin of Mexico can be characterized with a very high degree of confidence as more primate than Monte Albán IIIA in the Valley of Oaxaca, because the Middle Horizon rank-size curve presented by Blanton [2 (p. 196)] lies far outside the 90% confidence zone for the Monte Albán IIIA curve. Comparing the Middle Horizon to two earlier, more primate, patterns in the Valley of Oaxaca, we see that the difference between the Middle Horizon and Monte Albán Early I also has considerable significance. While our confidence is somewhat lower for a third comparison, we would still be more than 90% confident that even the extremely primate Formative period San José phase settlement system is not as primate as the Middle Horizon system, as

represented by Blanton's [2 (p. 196)] rank-size plot. (Since the Basin of Mexico survey was not yet complete when this rank-size plot was produced, it would be worth making the comparison anew with complete data.) Such confidence in this difference between regions would make it even more clear than before that efforts to understand the extremely primate character of the Teotihuacan settlement system are not misguided. We hope this final example helps show how the approach presented here makes rank-size analysis more powerful and more broadly applicable to study of the evolution of past settlement systems in comparative perspective.

Acknowledgements

Our thanks go to several colleagues who assisted us in the course of this research. Stephen Savage kindly consulted about the use of his program for evaluating the significance of rank-size patterns. Stephen Kowalewski and David Wilson generously provided settlement data for the Valley of Oaxaca and the Santa Valley, respectively, in electronic form. Gregory Indrisano read the paper in draft version and provided comments, and George Cowgill made particularly useful suggestions for clarification as a reviewer for this journal.

References

[1] M.J. Baxter, Methodological issues in the study of assemblage diversity, American Antiquity 66 (2001) 715–725.
[2] R.E. Blanton, The role of symbiosis in adaptation and sociocultural change, in: E.R. Wolf (Ed.), The Valley of Mexico: Studies in Prehispanic Ecology and Society, University of New Mexico Press, Albuquerque, 1976, pp. 181–201.
[3] R.E. Blanton, S. Kowalewski, G. Feinman, J. Appel, Monte Albán's Hinterland, Part I: The Prehispanic Settlement Patterns of the Central and Southern Parts of the Valley of Oaxaca, Mexico, Memoirs of the Museum of Anthropology, University of Michigan, No. 15, 1982.
[4] G.L. Cowgill, A selection of samplers: comments on archaeo-statistics, in: J.W. Mueller (Ed.), Sampling in Archaeology, University of Arizona Press, Tucson, 1975, pp. 258–293.
[5] G.L. Cowgill, The trouble with significance tests and what we can do about it, American Antiquity 42 (1977) 350–368.
[6] C.L. Crumley, Towards a locational definition of state systems of settlement, American Anthropologist 78 (1976) 59–73.
[7] R.D. Drennan, Statistics for Archaeologists: A Commonsense Approach, Plenum Press, New York, 1996.
[8] G.A. Johnson, Local Exchange and Early State Development in Southwestern Iran, Anthropological Papers, Museum of Anthropology, University of Michigan, No. 51, 1973.
[9] G.A. Johnson, Aspects of regional analysis in archaeology, Annual Review of Anthropology 6 (1977) 479–508.
[10] G.A. Johnson, Rank-size convexity and system integration: a view from archaeology, Economic Geography 56 (1980) 234–247.
[11] G.A. Johnson, Monitoring complex system integration and boundary phenomena with settlement size data, in: S.E. van der Leeuw (Ed.), Archaeological Approaches to the Study of Complexity, Universiteit van Amsterdam, Amsterdam, 1981, pp. 144–188.
[12] D. Kaufman, Measuring archaeological diversity: an application of the jackknife technique, American Antiquity 63 (1998) 73–85.
[13] S.A. Kowalewski, The evolution of primate regional systems, Comparative Urban Research 9 (1982) 60–78.
[14] S.A. Kowalewski, G.M. Feinman, L. Finsten, R.E. Blanton, L.M. Nicholas, Monte Albán's Hinterland, Part II: Prehispanic Settlement Patterns in Tlacolula, Etla, and Ocotlán, the Valley of Oaxaca, Mexico, Memoirs of the University of Michigan Museum of Anthropology, No. 23, 1989.
[15] L. Liu, Development of chiefdom societies in the middle and lower Yellow River valley in Neolithic China: a study of the Longshan Culture from the perspective of settlement patterns, Ph.D. Dissertation, Department of Anthropology, Harvard University (1994).
[16] L. Liu, Settlement patterns, chiefdom variability, and the development of early states in north China, Journal of Anthropological Archaeology 15 (1996) 237–288.
[17] C.Z. Mooney, R.D. Duvall, Bootstrapping: A Nonparametric Approach to Statistical Inference, Sage Publications, Newbury Park, CA, 1993.
[18] R.W. Paynter, Models of Spatial Inequality: Settlement Patterns in Historical Archaeology, Academic Press, New York, 1982.
[19] R.W. Paynter, Expanding the scope of settlement analysis, in: J.A. Moore, A.S. Keene (Eds.), Archaeological Hammers and Theories, Academic Press, New York, 1983, pp. 233–275.
[20] C.E. Pearson, Rank-size distributions and the analysis of prehistoric settlement systems, Journal of Anthropological Research 36 (1980) 453–462.
[21] S.H. Savage, Assessing departures from log-normality in the rank-size rule, Journal of Archaeological Science 24 (1997) 233–244.
[22] S.H. Savage, The RankSize program, v 2.03: guide to operation and interpretation, <http://archaeology.asu.edu/jordan/download.html>
[23] D.J. Wilson, The origins and development of complex prehistoric society in the lower Santa Valley, Peru: implications for theories of state origins, Journal of Anthropological Archaeology 2 (1983) 209–276.
[24] D.J. Wilson, Prehispanic Settlement Patterns in the Lower Santa Valley, Peru: A Regional Perspective on the Origins and Development of Complex North Coast Society, Smithsonian Institution Press, Washington, DC, 1988.
[25] G.K. Zipf, Human Behavior and the Principle of Least Effort, Harvard University Press, Cambridge, 1949.